arXiv:1701.04926v3 [cs.IT] 24 Feb 2017

Agglomerative Info-Clustering

Chung Chan, Ali Al-Bashabsheh and Qiaoqiao Zhou

***Abstract*—An agglomerative clustering of random variables is proposed, where clusters of random variables sharing the maximum amount of multivariate mutual information are merged successively to form larger clusters. Compared to the previous info-clustering algorithms, the agglomerative approach allows the computation to stop earlier when clusters of desired size and accuracy are obtained. An efficient algorithm is also derived based on the submodularity of entropy and the duality between the principal sequence of partitions and the principal sequence for submodular functions.**

I. Introduction

We consider the info-clustering paradigm proposed in [1]. It is a hierarchical clustering of a finite set of random variables (RVs) based on the multivariate mutual information (MMI) defined in [2]. The MMI is a natural extension of Shannon's mutual information to the multivariate case involving possibly more than two RVs. It was first proposed in [3] as a measure of mutual information after identifying the divergence upper bound of the secret key agreement problem [4] to be loose in the case with helper but tight in the no-helper case, which established an operational meaning of the MMI as the secrecy capacity [5]. The MMI was also shown to be equal to the undirected network coding throughput [6] under the matroidal undirected network link model [7].

The info-clustering solution was shown in [1] to coincide with an elegant mathematical structure called the principle sequence of partitions (PSP) [8] of a submodular function, namely, that of the entropy function [9] of the RVs to be clustered. This leads to an algorithm [1, Algorithm 3] similar to [10] that computes the clustering solution in $O(m^2\,\mathrm{SFM}(m))$ time, where $\mathrm{SFM}(m)$ is the time required to minimize a submodular function on a ground set of size m. In practice, however; 1) one may want to obtain clusters of a desired size rather than the entire hierarchy of clusters of different sizes; and 2) the entropy function needs to be estimated from data, which can be difficult for a large set of random variables [11]. These practical considerations motivate the search for an iterative info-clustering algorithm. A divisive clustering approach was proposed in [1, Algorithms 1 and 2] that breaks down the computation by splitting the entire set of RVs successively into increasingly smaller clusters. However, doing so appears to be inefficient, requiring $\Omega(m^3\,\mathrm{SFM}(m))$ time in the worst case. Furthermore, it computes the larger clusters first, the entropy function of which is more difficult to estimate from data, and so the error may be carried forward to subsequent computations of smaller clusters.

In this work, we propose an agglomerative info-clustering approach that aims to resolve the above issues. The idea is to start with smaller clusters first and merge them to form larger clusters successively. It turns out that the algorithm can be implemented more efficiently than the divisive approach, by relating the PSP to another structure called the principal sequence (PS) [12–14]. A similar duality between the PSP and PS was also used in [15] to relate info-clustering to the problem of feature selection. The contribution of this work is the derivation of a rigorous information-theoretic interpretation useful for the clustering problem, based on which further heuristics for estimation, approximation, or model reduction as in [1] can be developed.

C. Chan (email: cchan@inc.cuhk.edu.hk), A. Al-Bashabsheh, and Q. Zhou are with the Institute of Network Coding at the Chinese University of Hong Kong, the Shenzhen Key Laboratory of Network Coding Key Technology and Application, China, and the Shenzhen Research Institute of the Chinese University of Hong Kong.

The work described in this paper was supported by a grant from University Grants Committee of the Hong Kong Special Administrative Region, China (Project No. AoE/E-02/08), and supported partially by a grant from Shenzhen Science and Technology Innovation Committee (JSGG20160301170514984), the Chinese University of Hong Kong (Shenzhen), China.

The work of C. Chan was supported in part by The Vice-Chancellor's One-off Discretionary Fund of The Chinese University of Hong Kong (Project Nos. VCF2014030 and VCF2015007), and a grant from the University Grants Committee of the Hong Kong Special Administrative Region, China (Project No. 14200714).

II. Motivation

Before a rigorous and general treatment, we introduce the problem and results informally using the following example from [1, Figure 1a]: Consider the following RVs defined using the uniform and independent bits $\mathsf{X}_a, \mathsf{X}_b, \mathsf{X}_c$ and X_d

$$\begin{array}{lll} \mathsf{Z}_1 := (\mathsf{X}_a, \mathsf{X}_d), & \mathsf{Z}_2 := (\mathsf{X}_a, \mathsf{X}_d), & \mathsf{Z}_3 := \mathsf{X}_a, \\ \mathsf{Z}_4 := \mathsf{X}_b, & \mathsf{Z}_5 := \mathsf{X}_b, & \mathsf{Z}_6 := \mathsf{X}_c. \end{array} \tag{2.1}$$

Info-clustering [1] is a hierarchical clustering approach based on a multivariate measure of the information shared among (multiple) RVs. As will be discussed more precisely in a subsequent section, info-clustering provides clusters for different thresholds $\gamma \in \mathbb{R}$, where a cluster is an inclusion-wise maximal subset of RVs that share more than γ amount of information. (We do not regard a singleton as a cluster.) Since the correlation structure of the RVs in (2.1) is simple, let us for the moment define such a measure of information among Z_B, for any $B \subseteq \{1,\dots,6\}$ with $|B| \geq 2$, as the number of bits shared by Z_B and denote it by $I(\mathsf{Z}_B)$. Then we have

- $I(\mathsf{Z}_{\{1,\dots,6\}}) = 0$ since, e.g., $\mathsf{Z}_{\{1,\dots,5\}}$ and Z_6 share no bits. (Similarly, $I(\mathsf{Z}_{B\cup\{6\}}) = 0$ for $\emptyset \neq B \subseteq \{1,\cdots,5\}$; $I(\mathsf{Z}_{\{1,4\}}) = 0$; etc.)
- $I(\mathsf{Z}_{\{1,2,3\}}) = 1$ since Z_1, Z_2, and Z_3 share the bit X_a. (Similarly, $I(\mathsf{Z}_{\{1,3\}}) = 1 = I(\mathsf{Z}_{\{2,3\}})$.) We also have $I(\mathsf{Z}_{\{4,5\}}) = 1$ since Z_4 and Z_5 share the bit X_b.

- $I(\mathsf{Z}_{\{1,2\}}) = 2$ since Z_1 and Z_2 share the two bits X_a and X_d. (This is the only set of RVs sharing more than one bit.)
- There is no set of RVs that share more than two bits.

Hence, for all $\gamma \in \mathbb{R}$, the collection of clusters at threshold γ, denoted as $\mathcal{C}_\gamma(\mathsf{Z}_{\{1,\ldots,6\}})$, is given by [1, Figure 1b]

$$\mathcal{C}_\gamma(\mathsf{Z}_{\{1,\ldots,6\}}) = \begin{cases} \{\{1,\ldots,6\}\}, & \gamma < 0 \\ \{\{1,2,3\},\{4,5\}\}, & \gamma \in [0,1) \\ \{\{1,2\}\}, & \gamma \in [1,2) \\ \emptyset, & \gamma \geq 2. \end{cases} \tag{2.2}$$

For instance, for any $\gamma < 0$, there is only one cluster, namely, the entire set of RVs since any subset containing at least two RVs satisfies the threshold constraint and the entire set $\{1,\ldots,6\}$ is trivially the maximal one. For $\gamma = 0$, the threshold constraint dictates that we seek collections of RVs that share a strictly positive number of bits, i.e., one or two bits in this example. Of such sets, it is easy to verify that the maximal ones are $\{1,2,3\}$ and $\{4,5\}$. This remains to be the case for any $\gamma \in [0,1)$. For $\gamma = 1$, we seek collections of RVs that share more than one bits, i.e., two bits in this example. The set $\{1,2\}$ is the only such set, and so, it is trivially maximal. This remains to be the case for $\gamma \in [1,2)$. Finally, for $\gamma \geq 2$, we have no clusters since no set of RVs share more than two bits of information.

The reader may have observed the following hierarchical structure: Starting with the cluster $\{1,\ldots,6\}$ at sufficiently small threshold γ, the cluster breaks into two smaller clusters $\{1,2,3\}$ and $\{4,5\}$, where each is a maximal subset that has more shared information bits than the original cluster. Continuing in this fashion, the cluster $\{1,2,3\}$ breaks into the smaller cluster $\{1,2\}$, at which point no further breakage is possible and all clusters have been found. More generally, under a proper choice of the multivariate information measure, the hierarchical structure of the clusters persists, thereby allowing a divisive algorithm for finding the clusters [1, Algorithm 1]. The algorithm starts with the entire set of RVs then proceeds iteratively to break the clusters into smaller and smaller disjoint clusters.

In this work, we propose the reverse procedure for computing the clusters. Namely, an *agglomerative* approach where RVs gradually group into larger and larger clusters. In our example, starting with the singletons for sufficiently large value of γ, merge $\{1\}$ and $\{2\}$ into a cluster since such a merging results in the maximal subset $\{1,2\}$ with the maximum number of shared bits. Continue to merge $\{1,2\}$ with $\{3\}$ and merge $\{4\}$ with $\{5\}$ to form the maximal subsets $\{1,2,3\}$ and $\{4,5\}$ with the second largest number of shared bits. Continue in the same way to merge $\{1,2,3\}$, $\{4,5\}$, and $\{6\}$ into the cluster $\{1,\ldots,6\}$, at which point no further merging is possible and all clusters have been found.

The practicality of an agglomerative approach in comparison to a divisive one is that the former identifies clusters of larger amount of shared bits first. A larger amount of shared bits can be estimated from data more accurately, and so it serves as a stronger base for the subsequent estimation for finding clusters of smaller amount of shared bits. Moreover, as we will subsequently show, when the information measure is chosen to be the MMI in [2], the agglomerative approach is computationally more efficient compared to the divisive one.

III. Problem formulation

Given a random vector $\mathsf{Z}_V := (\mathsf{Z}_i \mid i \in V)$ where V is an ordered finite set of $|V| > 1$ RVs. The set of clusters at any threshold $\gamma \in \mathbb{R}$ is defined in [1] as

$$\mathcal{C}_\gamma(\mathsf{Z}_V) := \text{maximal}\{B \subseteq V \mid |B| > 1, I(\mathsf{Z}_B) > \gamma\}, \tag{3.1}$$

where $\text{maximal}\,\mathcal{F} := \{B \in \mathcal{F} \mid \nexists B' \in \mathcal{F}, B \subsetneq B'\}$ denotes the collection of inclusion-wise maximal elements of any collection $\mathcal{F}$ of subsets, and $I(\mathsf{Z}_B)$ is a multivariate information measure satisfying

$$I(\mathsf{Z}_{B_1 \cup B_2}) \geq \min\{I(\mathsf{Z}_{B_1}), I(\mathsf{Z}_{B_2})\} \tag{3.2}$$

for all $B_i \subseteq V$ with $|B_i| > 1$, $i \in \{1,2\}$ and $B_1 \cap B_2 \neq \emptyset$. It was shown in [1, Theorem 3] that the clusters form a laminar family, i.e., for any $\gamma' \leq \gamma''$, $C' \in \mathcal{C}_{\gamma'}(\mathsf{Z}_V)$, and $C'' \in \mathcal{C}_{\gamma''}(\mathsf{Z}_V)$, we have

$$C' \cap C'' = \emptyset \quad \text{or} \quad C' \supseteq C''. \tag{3.3}$$

In particular, clusters at the same threshold must be disjoint. Consequently, the clustering solution can be characterized as follows by set partitions of V, the collection of which is denoted as $\Pi(V)$.

Proposition 3.1 ([1, Theorems 1 and 4]) *The property* (3.2) *implies that the clustering solution* (3.1) *satisfies*

$$\mathcal{C}_\gamma(\mathsf{Z}_V) = \mathcal{P}_\ell \backslash \{\{i\} \mid i \in V\} \quad \text{for } \gamma \in [\gamma_\ell, \gamma_{\ell+1}], 0 \leq \ell \leq N$$

where N is a positive integer; $\gamma_0 := -\infty$ and $\gamma_{N+1} := \infty$ for convenience;

$$-\infty < \gamma_1 < \cdots < \gamma_N < \infty \tag{3.4}$$

is a sequence of distinct critical values from $\mathbb{R}$ (consisting of the thresholds at which the set of clusters changes); and

$$\mathcal{P}_0 = \{V\} \succ \mathcal{P}_1 \succ \cdots \succ \mathcal{P}_{N-1} \succ \mathcal{P}_N = \{\{i\} : i \in V\} \tag{3.5}$$

is a sequence of increasingly finer partitions of V from $\Pi(V)$, where $\mathcal{P}' \succeq \mathcal{P}$ means that

$$\forall C \in \mathcal{P}, \exists C' \in \mathcal{P}' : C \subseteq C', \tag{3.6}$$

and "$\succ$" denotes the strict inequality (i.e., the inclusion above is strict for at least one $C \in \mathcal{P}$). □

With the laminar structure, a divisive clustering algorithm was given in [1] to compute γ_ℓ and $\mathcal{P}_\ell$ by iteratively producing finer partitions from coarser ones, i.e., from $\ell = 1$ to $\ell = N$. The work in [2] considered a particularly meaningful multivariate information measure, namely, the *multivariate mutual information* (MMI) defined as

$$I(\mathsf{Z}_V) := \min_{\mathcal{P} \in \Pi'(V)} I_{\mathcal{P}}(\mathsf{Z}_V), \quad \text{where} \tag{3.7a}$$

$$I_{\mathcal{P}}(\mathsf{Z}_V) := \frac{1}{|\mathcal{P}| - 1}\left[\sum\nolimits_{C \in \mathcal{P}} H(\mathsf{Z}_C) - H(\mathsf{Z}_V)\right] \tag{3.7b}$$

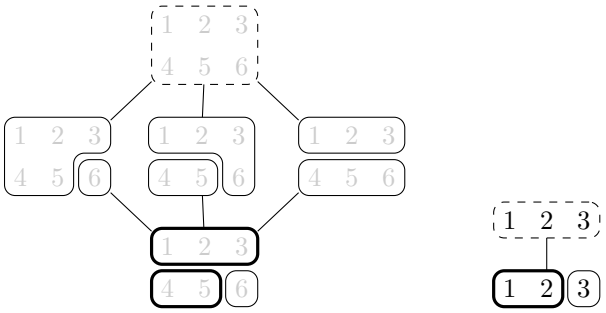

(a) (b)

Fig. 1: Optimal partitions of: (a) $I(\mathsf{Z}_V)$ and (b) $I(\mathsf{Z}_{\{1,2,3\}})$. In each case, the fundamental partition is the one at the bottom, where the associated clusters are circled with thick lines.

and $\Pi'(V) := \Pi(V) \setminus \{\{V\}\}$ is the set of all partitions of V into at least two non-empty disjoint subsets. The MMI was first proposed as a measure of mutual information in [3] and was later shown in [2] to satisfy (3.2), together with various other properties that naturally extend those of Shannon's mutual information to the multivariate case. The MMI was also later referred to in [16] as "shared information." An important result that inspired the info-clustering paradigm [1] is:

Proposition 3.2 ([2, Theorems 5.2 and 5.3]) *The optimal partitions achieving the MMI in (3.7) together with the trivial partition $\{V\}$ form a lattice w.r.t. (3.6). The minimum/finest optimal partition, denoted by $\mathcal{P}^*(\mathsf{Z}_V)$, satisfies*

$$\mathcal{P}^*(\mathsf{Z}_V)\backslash\{\{i\} \mid i \in V\} = \mathcal{C}_{I(\mathsf{Z}_V)}(\mathsf{Z}_V),$$

where $\mathcal{C}_\gamma$ is defined in (3.1) using the MMI. In particular, $\gamma_1 = I(\mathsf{Z}_V)$ and $\mathcal{P}_1 = \mathcal{P}^*(\mathsf{Z}_V)$ in Proposition 3.1. □

Together with the laminar structure in Proposition 3.1, it can be argued that any algorithm for computing $\mathcal{P}^*(\mathsf{Z}_V)$ can be applied iteratively for divisive info-clustering as in [1, Algorithm 2]. However, doing so appears to be less efficient compared to the alternative approach in [1, Algorithm 3] that computes all the clusters but not in any particular order. Nevertheless, in practice it is desirable to have an iterative approach that can stop when further computation is not of interest or is meaningless due to errors in estimating or approximating the entropies from data.

Example 1 As an illustration of Proposition 3.2 and the divisive approach, consider our motivating example. The finest optimal partition $\mathcal{P}^*(\mathsf{Z}_V) = \{\{1,2,3\},\{4,5\},\{6\}\}$ is shown in Fig. 1a. For completion, the figure also shows the lattice of optimal partitions stated in the proposition, where the trivial partition $\{V\}$ is indicated using a dashed line. Similarly, Fig. 1b shows the optimal partitions of $\mathsf{Z}_{\{1,2,3\}}$. Since $I(\mathsf{Z}_V) = 0$ and $I(\mathsf{Z}_{\{1,2,3\}}) = 1$, Proposition 3.2 asserts that $\mathcal{C}_0 = \{\{1,2,3\},\{4,5\}\}$ and $\mathcal{C}_1 = \{\{1,2\}\}$, in agreement with (2.2). Assuming an algorithm for computing the finest optimal partition, the divisive algorithm starts by computing $\mathcal{P}^*(\mathsf{Z}_V) = \{\{1,2,3\},\{4,5\},\{6\}\}$ (from which $I(\mathsf{Z}_V)$ is readily available), declares the non-singleton elements as clusters at threshold $I(\mathsf{Z}_V)$, and proceeds iteratively by picking any cluster (of size larger than two) and computing its finest optimal partition, etc. In our example, there is only one cluster, the set $\{1,2,3\}$ at threshold 0. The finest optimal partition of $\mathsf{Z}_{\{1,2,3\}}$ is shown in Fig. 1b, which results in the cluster $\{1,2\}$ at threshold $I(\mathsf{Z}_{\{1,2,3\}}) = 1$. After this, the divisive algorithm terminates in this example. □

Algorithm 1: Agglomerative info-clustering.

Data: Statistics of Z_V sufficient for calculating the entropy function $h(B)$ for $B \subseteq V := \{1,\dots,n\}$.

Result: The arrays L and PSP contain the γ_ℓ's and $\mathcal{P}_\ell$'s in Proposition 3.1. More precisely, for $1 \le \ell \le N$, the entries of the arrays PSP and L are $\mathsf{PSP}[s] = \mathcal{P}_\ell(\mathsf{Z}_V)$ and $\mathsf{L}[s] = \gamma_\ell(\mathsf{Z}_V)$ for $s = |\mathcal{P}_\ell(\mathsf{Z}_V)|$, and are null otherwise.

```
L, PSP ← empty arrays, each of length n;
PSP[n] ← {{i} : i ∈ V}, s ← n;
while s > 1 do
    (γ, P') ← Fuse(PSP[l]);
    L[s] ← γ;
    s ← |P'|;
    PSP[s] ← P';
end
```

Instead of the divisive approach, we consider here an agglomerative approach shown in Algorithm 1 that computes γ_ℓ's and $\mathcal{P}_\ell$'s iteratively from $\ell = N$ down to $\ell = 1$. We will give an efficient implementation of the subroutine `Fuse` in Algorithm 2 that computes γ_ℓ and $\mathcal{P}_{\ell-1}$ from $\mathcal{P}_\ell$ for any ℓ. In particular, we will show that it suffices to compute

$$I^*(\mathsf{Z}_V) := \max\{I(\mathsf{Z}_B) \mid B \subseteq V, |B| > 1\} \quad \text{and} \tag{3.8}$$

$$\mathcal{C}^*(\mathsf{Z}_V) := \text{maximal}\{B \subseteq V \mid |B| > 1, \; I(\mathsf{Z}_B) = I^*(\mathsf{Z}_V)\}, \tag{3.9}$$

which are clearly the last critical value γ_N and the non-singleton elements of the second last partition $\mathcal{P}_{N-1}$ respectively.

IV. Preliminaries

The ability to compute info-clustering solution efficiently stems from the submodularity of entropy [17], or equivalently the fact that mutual information is non-negative [18]. More precisely, by denoting the entropy function of Z_V as

$$h(B) := H(\mathsf{Z}_B) \quad \text{for } B \subseteq V, \tag{4.1}$$

(where the dependency on Z_V is implicit for convenience,) *submodularity* of h means that

$$h(B_1) + h(B_2) \ge h(B_1 \cup B_2) + h(B_1 \cap B_2) \tag{4.2}$$

for all $B_1, B_2 \subseteq V$. h is also said to be *normalized* as $h(\emptyset) = 0$, and non-decreasing as $h(B') \leq h(B)$ whenever $B' \subseteq B \subseteq V$. In combinatorial optimization [19], submodularity is well-known to give rise to polynomial-time solutions. The info-clustering problem, in particular, relies on the following closely related polynomial-time solvable structures.

A. Principal sequence of partitions

For any real number $\gamma \in \mathbb{R}$, define the residual entropy function [2] of a random vector Z_V as

$$h_\gamma(B) := h(B) - \gamma \qquad \text{for } B \subseteq V. \tag{4.3}$$

The residual entropy function is also submodular and its Dilworth truncation evaluated at V is defined as [19]

$$\hat{h}_\gamma(V) := \min_{\mathcal{P} \in \Pi(V)} h_\gamma[\mathcal{P}] \quad \text{where} \tag{4.4a}$$

$$h_\gamma[\mathcal{P}] := \sum\nolimits_{C \in \mathcal{P}} h_\gamma(C). \tag{4.4b}$$

Proposition 4.1 ([8]) *Submodularity (4.2) of h (4.1) implies that the set of optimal partitions to the Dilworth truncation (4.4a) forms a lattice, called the Dilworth truncation lattice, with respect to the partial order (3.6). Furthermore, if $\mathcal{P}'$ and $\mathcal{P}''$ are the optimal partitions for γ' and γ'' respectively, then $\gamma' < \gamma''$ implies $\mathcal{P}' \succeq \mathcal{P}''$.* □

In particular, the minimum/finest optimal partition exists and characterizes the info-clustering solution as follows:

Proposition 4.2 ([1, Corollary 2]) *For a finite set V with size $|V| > 1$ and a random vector Z_V,*

$$\mathcal{C}_\gamma(\mathsf{Z}_V) = \Big[\min\{\mathcal{P} \in \Pi(V) \mid h_\gamma[\mathcal{P}] = \hat{h}_\gamma(V)\}\Big] \setminus \{\{i\} \mid i \in V\}, \tag{4.5}$$

namely, the non-singleton elements of the finest optimal partition to the Dilworth truncation (4.4a). □

Hence, in Proposition 3.1, the sequence of $\mathcal{P}_\ell$ for ℓ from 1 to N with the corresponding γ_ℓ also characterizes the minimum optimal partitions to (4.4a) for all $\gamma \in \mathbb{R}$, and is known as the principal sequence of partitions (PSP) of h, introduced in [8].

Example 2 For the motivating example, the Dillworth truncation $\hat{h}_\gamma(V)$ (4.4a) is shown in Fig. 2a. The Dillowrth truncation is piecewise linear with at most $|V|$ turning points. A turning point $p_i := (\gamma_i, \hat{h}_{\gamma_i}(V))$ occurs when (4.4a) has more than one solution. The collection of such optimal solutions is the Dillowrth truncation lattice at γ_i indicated in Proposition 4.1. (The Dillowrth truncation lattice is not shown in Fig. 2.) The finest optimal partition at γ_i remains (while all other partitions seieze to be) optimal until the next turning point. In other words, the finest optimal partition at γ_i determines the line segment that follows γ_i. The sequence of the finest optimal partitions is the PSP, whch is shown in Fig. 2b. □

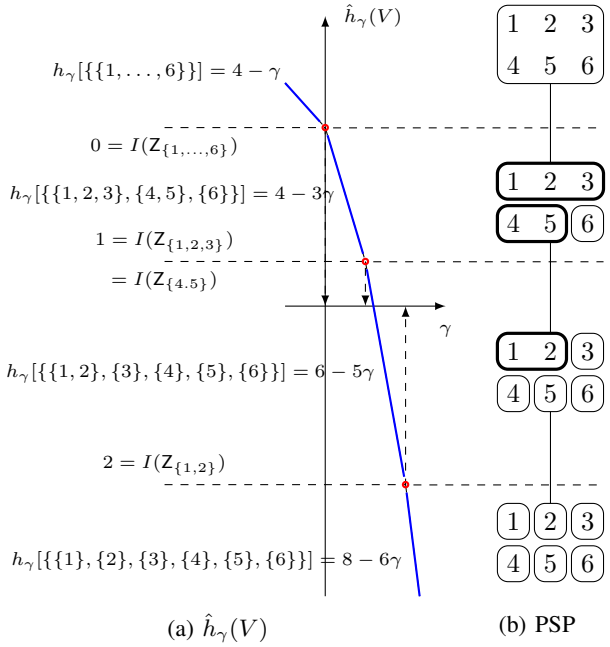

(a) $\hat{h}_\gamma(V)$ (b) PSP

Fig. 2: Dilworth truncation

B. Principal sequence and minimum norm base

Consider any submodular function $f : 2^U \to \mathbb{R}$ (4.2) on the finite ground set U. For $\lambda \in \mathbb{R}$, define

$$S_\lambda(f) := \max \arg\min_{B \subseteq V} f(B) - \lambda|B|. \tag{4.6}$$

where max is the inclusion-wise maximum. Note that the minimization in (4.6) is a *submodular function minimization (SFM)* since the function $B \subseteq U \mapsto f(B) - \lambda|B|$ is also submodular. It is well-known that the minimizers form a lattice with respect to set inclusion [19], and so the maximum in (4.6) exists and is unique.

Proposition 4.3 ([12, 14]) *$S_\lambda(f)$ for $\lambda \in \mathbb{R}$ satisfies*

$$S_{\lambda'}(f) \subseteq S_{\lambda''}(f) \quad \textit{iff} \quad \lambda' \leq \lambda''$$

and is referred to as the principal sequence (PS). □

Without loss of generality, we assume f is normalized, i.e., $f(\emptyset) = 0$, because we can redefine f as $f - f(\emptyset)$ without affecting the solutions to the SFM in (4.6), i.e. the PS is invariant to constant shift in the submodular function. The polyhedron $\mathrm{P}(f)$ and base polyhedron $\mathrm{B}(f)$ of the submodular function f are defined as

$$\mathrm{P}(f) := \{x_U \in \mathbb{R}^U \mid x(B) \leq f(B), \forall B \subseteq U\} \tag{4.7}$$

$$\mathrm{B}(f) := \{x_U \in \mathrm{P}(f) \mid x(U) = f(U)\} \tag{4.8}$$

where $x_U := (x_i \mid i \in U)$ and $x(B) := \sum_{i \in B} x_i$ for convenience. ($\mathrm{P}(f)$ and $\mathrm{B}(f)$ are non-empty as $f(\emptyset) \geq 0$.) With $\|x_U\|$ denoting the Euclidean norm of the vector x_U, the following holds.

Proposition 4.4 ([12, 20]) *For any normalized submodular function f (4.2),*

$$\min\{\|x_U\| \mid x_U \in \mathrm{B}(f)\} \tag{4.9}$$

has a unique solution x^*_U, called the minimum (Euclidean) norm base, which satisfies

$$x^*_i(f) = \min\{\lambda \in \mathbb{R} \mid i \in S_\lambda(f)\} \quad \forall i \in U, \text{or equiv.,} \tag{4.10a}$$
$$S_\lambda(f) = \{i \in U \mid x^*_i(f) \le \lambda\}, \quad \forall \lambda \in \mathbb{R}, \tag{4.10b}$$

where the equivalence follows directly from Proposition 4.3. □

The minimum norm base may be computed using Wolfe's minimum norm point algorithm as in [20], and so as the PS by (4.10b). Conversely, by (4.10a), the minimum norm base can also be computed by any SFM algorithm that solves S_λ for any λ, but the minimum norm point algorithm was shown [20] empirically to perform well compared to other submodular function minimization algorithms.

V. Main results

Algorithm 2: Implementation of `Fuse` in Algorithm 1.

Data: $\mathcal{P}$ is equal to $\mathcal{P}_\ell$ for some $1 \le \ell \le N$ in Proposition 3.1.
Result: $(\gamma, \mathcal{P}')$ is equal to $(\gamma_\ell, \mathcal{P}_{\ell-1})$.

```
Enumerate P as {C_1, ..., C_k} for some k > 1 and disjoint C_i's;
x ← empty array of size k;
for j = 1 to k do
    x[j] = MinNormBase (B ↦ h(⋃_{i∈B∪{j}} C_i) − Σ_{i∈B∪{j}} h(C_i), {j+1, ..., k});
end
γ ← − min_{i,j:1≤i<j≤k} x[j][i],   P' ← ∅;
for j = 1 to k do
    if C_j ∉ ⋃P' then
        add {C_j} ∪ {C_i | i∈{j+1, ..., k}, x[j][i]≤−γ} to P';
    end
end
function MinNormBase (f, U):
    return an array x (indexed by U) that solves (4.9).
end
```

The main result is the implementation of `Fuse` in Algorithm 2 that computes the PSP and therefore the info-clustering solution iteratively from the finer partitions to the coarser ones. This is done by computing the PS using a subroutine `MinNormBase` that computes the minimum norm base that solves (4.9). An explicit implementation of this subroutine can be found in [20] using Wolfe's minimum norm point algorithm. The current abstraction also allows further approximations or simplifications for special source models as in [1].

To explain Algorithm 2, we first simplify what `Fuse` should compute as follows.

Theorem 5.1 *Consider $\mathcal{P}_\ell$ and γ_ℓ defined as in Proposition 3.1 for $1 \le \ell \le N$ and write*

$$\mathcal{P}_\ell = \{C_j \mid j \in U\} \quad \text{and} \quad \mathsf{Z}'_j := \mathsf{Z}_{C_j}$$

for some index set U and disjoint subset C_j for $j \in U$. Then, for $1 \le \ell \le N$ we have

$$\gamma_\ell = \max_{\mathcal{F} \subseteq \mathcal{P}_\ell, |\mathcal{F}|>1} I(\mathsf{Z}_{\bigcup \mathcal{F}}) \tag{5.1a}$$
$$= I^*(\mathsf{Z}'_U) \tag{5.1b}$$
$$\mathcal{P}_{\ell-1} \setminus \mathcal{P}_\ell = \text{maximal}\left\{ \bigcup \mathcal{F} \,\middle|\, \mathcal{F} \in \arg\max_{\mathcal{F} \subseteq \mathcal{P}_\ell : |\mathcal{F}|>1} I(\mathsf{Z}_{\bigcup \mathcal{F}}) \right\} \tag{5.2a}$$
$$= \left\{ \bigcup_{j \in B} C_j \mid B \in \mathcal{C}^*(\mathsf{Z}'_U) \right\}, \tag{5.2b}$$

where $\bigcup \mathcal{F} := \bigcup_{B \in \mathcal{F}} B$ for convenience. More precisely, (5.1a) and (5.2a) follow more generally from the property (3.2), while (5.1b) and (5.2b) follow from the definition (3.7) of the MMI. □

Proof See Appendix A. ■

It follows from (5.1b) and (5.2b) that it suffices to compute I^* (3.8) and $\mathcal{C}^*$ (3.9). This can be done using a minimum norm base algorithm as follows. Define

$$J_{\mathrm{T}}(\mathsf{Z}_V) := I_{\{\{i\} | i \in V\}}(\mathsf{Z}_V) = \frac{1}{|V|-1}\left[\sum_{i \in V} H(\mathsf{Z}_i) - H(\mathsf{Z}_V)\right], \tag{5.3}$$

which is called the normalized total correlation [2] as it is the same as Watanabe's total correlation [21] except for the additional normalization factor of $\frac{1}{|V|-1}$.

Theorem 5.2 *With entropy function h (4.1) for Z_V, define*

$$U_j := \{i \in V \mid i > j\} \quad \text{for } j \in V, \text{ assuming } V \text{ is ordered.}$$
$$g_j(B) := h(B \cup \{j\}) - \sum\nolimits_{i \in B \cup \{j\}} h(\{i\}) \quad \text{for } B \subseteq U_j,$$

which is a normalized submodular function. Then,

$$I^*(\mathsf{Z}_V) = \max_{C \subseteq V : |C|>1} J_T(\mathsf{Z}_C) \tag{5.4a}$$
$$= -\min_{j \in V} \min_{i \in U_j} x^{(j)}_i \tag{5.4b}$$
$$\mathcal{C}^*(\mathsf{Z}_V) = \text{maximal} \arg\max_{C \subseteq V : |C|>1} J_T(\mathsf{Z}_C) \tag{5.5a}$$
$$= \text{maximal}\left\{ \{j^*\} \cup \arg\min_{i \in U_{j^*}} x^{(j^*)}_i \,\middle|\, j^* \in \arg\min_{j \in V} \left(\min_{i \in U_j} x^{(j)}_i \right) \right\}, \tag{5.5b}$$

where $x^{(j)}_{U_j}$ is the minimum norm base for g_j (4.9). □

Proof See Appendix A. ■

(5.4a) and (5.5a) essentially eliminate the need for minimization over partitions in calculating the MMI in (3.8)

and (3.9). They serve as an intermediate step that leads to (5.4b) and (5.5b), which relate the last critical value (3.8) and the second last partition (3.9) to the minimum norm base. Together with Proposition 4.4, we have the complete implementation of `Fuse` as shown in Algorithm 2 using a minimum norm base algorithm.

Example 3 As an illustration of Theorem 5.2 (and the agglomerative algorithm), consider our running example with $V = \{1, \ldots, 6\}$. The minimum norm base $x^{(j)}_{U_j}$ of g_j is given as

$$x^{(j)}_{U_j} = \begin{cases} (\quad \underset{\cdots}{-2}, & -1, & -0.5, & -0.5, & 0), & j = 1 \\ (\quad\quad\quad & -1, & -0.5, & -0.5, & 0), & j = 2 \\ (\quad\quad\quad & & -0.5, & -0.5, & 0), & j = 3 \\ (\quad\quad\quad & & & -1, & 0), & j = 4 \\ (\quad\quad\quad & & & & 0), & j = 5. \end{cases} \qquad \begin{matrix} 1 & 2 & 3 & 4 & 5 & 6 \end{matrix}$$

By (5.4b), we have $I^*(\mathsf{Z}_V) = 2$ and by (5.5b), we have $\mathcal{C}^*(\mathsf{Z}_V) = \{\{1, 2\}\}$. In other words, starting with the partition into singletons in the agglomerative algorithm, the function `Fuse` returns the threshold value 2 and the partition $\{\{1, 2\}, \{3\}, \{4\}, \{5\}, \{6\}\}$. Let $\mathsf{Z}'_1 = \mathsf{Z}_{\{1,2\}}$ and $\mathsf{Z}'_i = \mathsf{Z}_{i+1}$ for $i = 2, \ldots, 5$. The minimum norm base $x^{(j)}_{U_j}$ of g_j (defined using the entropy function of $\mathsf{Z}'_{\{1,\ldots,5\}}$) is given as

$$x^{(j)}_{U_j} = \begin{cases} (\quad \underset{\cdots}{-1}, & -0.5, & -0.5, & 0), & j = 1 \\ (\quad\quad\quad & -0.5, & -0.5, & 0), & j = 2 \\ (\quad\quad\quad & & \underset{\cdots}{-1}, & 0), & j = 3 \\ (\quad\quad\quad & & & 0), & j = 4 \end{cases} \qquad \begin{matrix} 1 & 2 & 3 & 4 & 5 \end{matrix}$$

By (5.4b), we have $I^*(\mathsf{Z}'_{\{1,\ldots,5\}}) = 1$ and by (5.5b), we have $\mathcal{C}^*(\mathsf{Z}'_{\{1,\ldots,5\}}) = \{\{1, 2\}, \{3, 4\}\}$, which in terms of Z_V, results in the clusters $\{\{1, 2, 3\}, \{4, 5\}\}$. In other words, when called with the input partition $\{\{1, 2\}, \{3\}, \{4\}, \{5\}, \{6\}\}$, the function `Fuse` returns the threshold value 1 and the partition $\{\{1, 2, 3\}, \{4, 5\}, \{6\}\}$. Let $\mathsf{Z}''_1 = \mathsf{Z}_{\{1,2,3\}}$, $\mathsf{Z}''_2 = \mathsf{Z}_{\{4,5\}}$, and $\mathsf{Z}''_3 = \mathsf{Z}_6$. The minimum norm base $x^{(j)}_{U_j}$ of g_j (defined using the entropy function of $\mathsf{Z}''_{\{1,2,3\}}$) is given as

$$x^{(j)}_{U_j} = \begin{cases} (\quad \underset{\cdots}{0}, & \underset{\cdots}{0}), & j = 1 \\ (\quad\quad\quad & \underset{\cdots}{0}), & j = 2 \end{cases} \qquad \begin{matrix} 1 & 2 & 3 \end{matrix}$$

By (5.4b), we have $I^*(\mathsf{Z}''_{\{1,2,3\}}) = 0$ and by (5.5b), we have $\mathcal{C}^*(\mathsf{Z}''_{\{1,2,3\}}) = \{1, 2, 3\}$, which in terms of Z_V, results in the cluster $\{1, \ldots, 6\}$. In other words, when called with the input partition $\{\{1, 2, 3\}, \{4, 5\}, \{6\}\}$, the function `Fuse` returns the threshold value 0 and the trivial partition $\{1, \ldots, 6\}$, where at this point the agglomerative algorithm terminates. □

The complexity of the algorithm is mainly due to the computation of the minimum norm base in line 4. This computation is repeated at most $|V|$ times. With $\mathrm{MNP}(l)$ being the complexity of the minimum norm base algorithm for ground set of size l, then `Fuse` runs in time $O(|V|\,\mathrm{MNP}(|V|))$. Since the agglomerative info-clustering algorithm in Algorithm 1 invokes function `Fuse` $N - 1 \leq |V| - 1$ times, it runs in time $O(|V|^2\,\mathrm{MNP}(|V|))$, which is equivalent to that of [1, Algorithm 3], assuming that the submodular function minimization therein is implemented by the minimum norm base algorithm, i.e., with $\mathrm{SFM} = \mathrm{MNP}$. The divisive info-clustering algorithm in [1, Algorithm 2] makes $N-1$ calls to a subroutine that calculates the fundamental partition. However, computing the fundamental partition appears to take time $O(|V|^2\,\mathrm{MNP}(|V|))$, which would lead to an overall complexity of $O(|V|^3\,\mathrm{MNP}(|V|))$ for the divisible clustering. Hence, the agglomerative info-clustering appears more efficient.

VI. Conclusion

To address the concern of entropy estimation and computational complexity, we have proposed an agglomerative info-clustering approach that merges smaller clusters into larger clusters successively. To the best of our knowledge, this is the fastest info-clustering algorithm without any approximation. As mentioned in [1], however, faster algorithms are possible under special source models, such as the Markov tree model or Chow–Liu tree approximation [22]. For the graphical source models, the PSP can be computed more efficiently using a parametric maxflow algorithm as in [23]. The info-clustering algorithm can also be used to compute the solution of some related problems such as the optimal discussion rate tuple for successive omniscience [24].

Appendix A
Proofs of Main Results

Proof (Theorem 5.1) First, we prove (5.1a) and (5.2a) using the general property (3.2) instead of the precise definition (3.7).

We first argue that, for any feasible solution $\mathcal{F}$ to the r.h.s. of (5.1a),

$$I(\mathsf{Z}_{\bigcup \mathcal{F}}) \overset{\text{(a)}}{\leq} \gamma_\ell.$$

Suppose to the contrary that $I(\mathsf{Z}_{\bigcup \mathcal{F}}) > \gamma_\ell$. Then, there exists $C \supseteq \bigcup \mathcal{F}$ such that $C \in \mathcal{C}_{\gamma_\ell}(\mathsf{Z}_V)$ by the definition (3.1) of clusters. However, $\mathcal{C}_{\gamma_\ell}(\mathsf{Z}_V) \subseteq \mathcal{P}_\ell$ by Proposition 3.1, which contradicts the fact that $\mathcal{F} \subseteq \mathcal{P}_\ell$ with $|\mathcal{F}| > 1$.

Next, we show that (a) can be achieved with equality for some feasible solution $\mathcal{F}$. Consider any $C \in \mathcal{P}_{\ell-1} \setminus \mathcal{P}_\ell$. (Such a C exists since $\mathcal{P}_\ell$ is strictly finer than $\mathcal{P}_{\ell-1}$.) Then, we have

$$C \overset{\text{(b)}}{=} \bigcup \mathcal{F} \text{ for some } \mathcal{F} \subseteq \mathcal{P}_\ell : |\mathcal{F}| > 1,$$

i.e., for some feasible solution $\mathcal{F}$.

By Proposition 3.1, we have

$$C \in \mathcal{C}_\gamma(\mathsf{Z}_V) \quad \text{for all } \gamma \in [\gamma_{\ell-1}, \gamma_\ell),$$

and so $I(\mathsf{Z}_C) \geq \gamma_\ell$, i.e., larger than all values in the interval. The reverse inequality also holds by (a) and (b). Hence, we have

$$I(\mathsf{Z}_C) \overset{\text{(c)}}{=} \gamma_\ell,$$

which implies (5.1a) as desired.

Now, we argue that the above construction gives all the optimal solutions to the r.h.s. of (5.1a), hence establishing (5.2a). For any $C \in \mathcal{P}_{\ell-1} \setminus \mathcal{P}_\ell$, (b) and (c) implies that "$\subseteq$" holds for (5.2a), because the fact that $C \in \mathcal{C}_{\gamma_{\ell-1}}(\mathsf{Z}_V)$ (by Proposition 3.1) means that it is maximal by the definition (3.1) of clusters.

To argue the reverse inclusion "$\supseteq$", consider any $\mathcal{F}$ belonging to the r.h.s. of (5.2a). By (5.1a),

$$I(\mathsf{Z}_{\bigcup \mathcal{F}}) = \gamma_\ell > \gamma_{\ell-1}$$

and so $\bigcup \mathcal{F} \in \mathcal{C}_{\gamma_{\ell-1}}$ by the definition (3.1) of clusters and the maximality of $\bigcup \mathcal{F}$. This completes the poof of (5.2a).

Consider proving (5.1b) and (5.2b). For any optimal solution $\mathcal{F}$ to (5.1a), we have

$$\mathcal{P}^*(\mathsf{Z}_{\bigcup \mathcal{F}}) = \mathcal{F} = \{C_j \mid j \in B\}$$

for some $B \subseteq U$, where the first equality is by Proposition 3.2 since $I(\mathsf{Z}_C) > \gamma_\ell$ for all $C \in \mathcal{F}$ such that $C > 1$. Hence,

$$\begin{aligned} I(\mathsf{Z}_{\bigcup \mathcal{F}}) = I_{\mathcal{P}^*(\mathsf{Z}_{\bigcup \mathcal{F}})}(\mathsf{Z}_{\bigcup \mathcal{F}}) &= I_{\mathcal{F}}(\mathsf{Z}_{\bigcup \mathcal{F}}) \\ &= I_{\{\{j\} \mid j \in B\}}(\mathsf{Z}'_B) \geq I(\mathsf{Z}'_B) \end{aligned}$$

where the inequality follows from the fact that partition $\{\{j\} \mid j \in B\}$ into singletons may not be the optimal partition of B for $I(\mathsf{Z}'_B)$. The reverse inequality also holds because, for all $\mathcal{P}' \in \Pi'(B)$, define

$$\mathcal{P} := \left\{ \bigcup_{j \in C'} C_j \mid C' \in \mathcal{P}' \right\} \in \Pi'\Big(\bigcup \mathcal{F}\Big),$$

we have $I_{\mathcal{P}'}(\mathsf{Z}'_B) = I_{\mathcal{P}}(\mathsf{Z}_{\bigcup \mathcal{F}}) \geq I(\mathsf{Z}_{\bigcup \mathcal{F}})$. Here, the inequality follows from the fact that $\mathcal{P}$ may not be the optimal partition of $\mathsf{Z}_{\bigcup \mathcal{F}}$. This completes the proof of Theorem 5.1.■

PROOF (THEOREM 5.2) Consider $\gamma_N, \mathcal{P}_{N-1}$, and $\mathcal{P}_N$ as in Proposition 3.1 and let h_γ be as in (4.3). By Proposition 4.2, we have for all $C \in \mathcal{P}_{N-1} \setminus \mathcal{P}_N$ that

$$\begin{aligned} h_{\gamma_N}(C) &\overset{\text{(a)}}{=} \sum_{i \in C} h_{\gamma_N}(\{i\}) \quad \text{or equivalently,} \\ \gamma_N &\overset{\text{(b)}}{=} J_{\mathrm{T}}(\mathsf{Z}_C). \end{aligned}$$

- The equivalence between (a) and (b) follows immediately from the definition (5.3) of J_{T}. More precisely, (b) is equivalent to

$$\begin{aligned} \gamma_N &= \frac{1}{|C|-1}\left[\sum_{i \in C} h(\{i\}) - h(C)\right] && \text{by (4.1)} \\ \iff (|C|-1)\gamma_N &= \sum_{i \in C} h(\{i\}) - h(C) && \because |C| > 1 \\ \iff h_{\gamma_N}(C) &= \sum_{i \in C} h_{\gamma_N}(\{i\}) && \text{by (4.3)} \end{aligned}$$

  which is equivalent to (a) as desired.
- "$\geq$" for (a) follows from

$$h_{\gamma_N}\big[\{C\} \cup \{\{i\} \mid i \in V \setminus C\}\big] \geq h_{\gamma_N}(\mathcal{P}_N)$$

  since $\mathcal{P}_N$ is optimal to $\hat{h}_{\gamma_N}(V)$ (4.4) by construction.
- To explain "$\leq$" for (a), note that $\mathcal{P}_{N-1}$ is an optimal partition to the Dilworth truncation (4.4a) for $\gamma \in [\gamma_{N-1}, \gamma_N)$ by Proposition 4.2. By continuity of $h_\gamma[\mathcal{P}_{N-1}]$ (4.4b) with respect to γ, we have that $\mathcal{P}_{N-1}$ is also optimal for $\gamma = \gamma_N$, i.e.,

$$h_{\gamma_N}[\mathcal{P}_N] \overset{\text{(c)}}{=} h_{\gamma_N}[\mathcal{P}_{N-1}],$$

  by the optimality of $\mathcal{P}_N$. Hence, since

$$\sum_{i \in \mathcal{P}_N} h_{\gamma_N}(\{i\}) = \sum_{C \in \mathcal{P}_{N-1}} \sum_{i \in C} h_{\gamma_N}(\{i\}),$$

  we have,

$$\sum_{C \in \mathcal{P}_{N-1} \setminus \mathcal{P}_N} \left[h_{\gamma_N}(C) - \sum_{i \in C} h_{\gamma_N}(\{i\}) \right] = 0.$$

  Each term in the bracket is non-negative as argued before, and so must be equal to zero.

Consider any maximal solution C' to the r.h.s. of (5.4a). Then,

$$\begin{aligned} \gamma_N &\overset{\text{(d)}}{\leq} J_{\mathrm{T}}(\mathsf{Z}_{C'}) \quad \text{or equivalently,} \\ h_{\gamma_N}(C') &\overset{\text{(e)}}{\leq} \sum_{i \in C'} h_{\gamma_N}(\{i\}), \end{aligned}$$

which follow from (b) and (a) respectively because the (5.4a) does not require $C \in \mathcal{P}_{N-1} \setminus \mathcal{P}_N$. With

$$\mathcal{P}' := \{C'\} \cup \{\{i\} \mid i \in V \setminus C'\},$$

(e) implies

$$h_{\gamma_N}[\mathcal{P}'] \leq h_{\gamma_N}[\mathcal{P}_N],$$

which must hold with equality by the optimality of $\mathcal{P}_N$. Therefore, (d) also holds with equality, and so we have (5.4a) since $\gamma_N = I^*(\mathsf{Z}_V)$. Note that, by (c), $\mathcal{P}_{N-1}$ is also an optimal partition. Furthermore, it is the largest/coarsest such partition since it is optimal for some $\gamma < \gamma_N$ and therefore larger/coarser than all optimal partitions for $\gamma = \gamma_N$ by Proposition 4.1. Therefore, the set of maximal optimal solutions to the r.h.s. of (5.4a) is $\mathcal{P}_{N-1} \setminus \mathcal{P}_N$, which is $\mathcal{C}^*(\mathsf{Z}_V)$ trivially, and so we have (5.5a).

To prove (5.4b) and (5.5b), let $\gamma^* = I^*(\mathsf{Z}_V)$. Then, it follows from (5.4a) that

$$\begin{aligned} 0 &= \min_{C \subseteq V : |C| > 1} \gamma^* - J_{\mathrm{T}}(\mathsf{Z}_C) \\ &\overset{\text{(f)}}{=} \min_{C \subseteq V : |C| > 1} (|C|-1)\gamma^* + h(C) - \sum_{i \in C} h(\{i\}) \\ &\overset{\text{(g)}}{=} \min_{j \in V} \min_{B \subseteq U_j : |B| \geq 1} g_j(B) + \gamma^* |B|. \end{aligned}$$

- (f) is obtained by the definition (5.3) of J_{T} and multiplying both sides of the equality by $|C| - 1 > 0$, which does neither violate the equality nor change the set of solutions;

- (g) is obtained by changing the variable C to (j, B) using the bijection that sets

$$j := \min_{i \in C} i \quad \text{and} \quad B := C \setminus \{j\},$$

which is possible because $|C| > 1$. We have also applied $|B| = |C| - 1$ and the definition of g_j to rewrite the expression in the minimization.

Consider any optimal solution j^* to the R.H.S. of (g). Then, we have

$$S_\lambda(g_{j^*}) \begin{cases} \overset{\text{(h)}}{\neq} \emptyset & \lambda = -\gamma^* \\ \overset{\text{(i)}}{=} \emptyset & \lambda < -\gamma^*. \end{cases}$$

- (h) is because, by (g),

$$\begin{aligned} \emptyset &\neq \max \operatorname*{arg\,min}_{B \subseteq U_{j^*} : |B| \geq 1} g_{j^*}(B) + \gamma^* |B| \\ &= \max \operatorname*{arg\,min}_{B \subseteq U_{j^*}} g_{j^*}(B) + \gamma^* |B| \\ &= S_{-\gamma^*}(g_{j^*}) \end{aligned}$$

The last equality is by the definition (4.6) of $S_{-\gamma^*}$. The first equality is because allowing $B = \emptyset$ does not change the minimum value 0, since

$$g_{j^*}(\emptyset) + \gamma^* |\emptyset| = 0$$

as $g_{j^*}(\emptyset)$. Doing so also does not affect the maximum minimizer since the new optimal solution introduced, namely $\emptyset$, cannot be maximum trivially.

- To explain (i), note that (g) implies for all $B \subseteq U_{j^*} : |B| > 1$ that

$$\begin{aligned} g_{j^*}(B) + \gamma^* |B| &\geq 0 && \text{and so} \\ g_{j^*}(B) - \lambda |B| &\geq (-\lambda - \gamma^*)|B| && \forall \lambda \in \mathbb{R} \\ &> 0 && \forall \lambda < -\gamma^*. \end{aligned}$$

In other words, for $\lambda < -\gamma^*$, we have that $\emptyset$ is the unique solution to $\min_{B \subseteq U_{j^*}} g_{j^*}(B) + \gamma^* |B|$, which implies (i) by the definition (4.6) of S_λ.

Now, (h) and (i) implies that

$$\begin{aligned} -\gamma^* &= \sup\{\lambda \in \mathbb{R} \mid S_\lambda(g_{j^*}) = \emptyset\} \\ &= \min_{i \in U_{j^*}} \min_{\lambda \in \mathbb{R} : i \in S_\lambda(g_{j^*})} \lambda \\ &\overset{\text{(j)}}{=} \min_{i \in U_{j^*}} x_i^{(j^*)} \end{aligned}$$

where the last equality (j) is by (4.10a) since $x_{U_{j^*}}^{(j^*)}$ denotes the minimum norm base for g_{j^*}. The equality implies (5.4b) as desired.

To prove (5.5b), consider any set C^* that belongs to the r.h.s. of (5.5a). Applying the bijection

$$j^* := \min_{i \in C^*} i \quad \text{and} \quad B^* := C^* \setminus \{j^*\},$$

(j^*, B^*) is a solution to the r.h.s. of (g). By the inclusion-wise maximality of C^*, the set B^* is also a maximal (the maximum) solution, i.e.,

$$\begin{aligned} B^* &\overset{\text{(k)}}{=} S_{-\gamma^*}(g_{j^*}) \\ &= \{i \in U_{j^*} \mid i \in S_{-\gamma^*}(g_{j^*})\} \\ &\overset{\text{(l)}}{=} \{i \in U_{j^*} \mid x_i^{(j^*)} \leq -\gamma^*\} \\ &\overset{\text{(m)}}{=} \operatorname*{arg\,min}_{i \in U_{j^*}} x_i^{(j^*)}, \end{aligned}$$

where (k) is by (4.6); (l) is by (4.10b); and (m) is by (j). Hence, C^* also belongs to the r.h.s. of (5.5b).

Conversely, if (j^*, B^*) is an optimal solution to the r.h.s. of (g), it can be argued easily that $C^* := \{j\}^* \cup B^*$ is also a solution to the r.h.s. of (5.4a). The maximal such C^* therefore belongs to the l.h.s. of (5.5b) as desired. This completes the proof. ∎

REFERENCES

[1] C. Chan, A. Al-Bashabsheh, Q. Zhou, T. Kaced, and T. Liu, "Info-clustering: A mathematical theory for data clustering," *IEEE Transactions on Molecular, Biological and Multi-Scale Communications*, vol. 2, no. 1, pp. 64–91, June 2016.

[2] C. Chan, A. Al-Bashabsheh, J. Ebrahimi, T. Kaced, and T. Liu, "Multivariate mutual information inspired by secret-key agreement," *Proceedings of the IEEE*, vol. 103, no. 10, pp. 1883–1913, Oct 2015.

[3] C. Chan, "On tightness of mutual dependence upperbound for secret-key capacity of multiple terminals," *arXiv preprint arXiv:0805.3200*, 2008.

[4] I. Csiszár and P. Narayan, "Secrecy capacities for multiple terminals," *IEEE Trans. Inf. Theory*, vol. 50, no. 12, Dec. 2004.

[5] C. Chan and L. Zheng, "Mutual dependence for secret key agreement," in *Proceedings of 44th Annual Conference on Information Sciences and Systems*, 2010.

[6] C. Chan, "The hidden flow of information," in *Proc. IEEE Int. Symp. on Inf. Theory*, St. Petersburg, Russia, Jul. 2011.

[7] ——, "Matroidal undirected network," in *Information Theory Proceedings (ISIT), 2012 IEEE International Symposium on*, July 2012, pp. 1498–1502.

[8] H. Narayanan, "The principal lattice of partitions of a submodular function," *Linear Algebra and its Applications*, vol. 144, no. 0, pp. 179 – 216, 1990.

[9] R. W. Yeung, *Information Theory and Network Coding*. Springer, 2008.

[10] K. Nagano, Y. Kawahara, and S. Iwata, "Minimum average cost clustering." in *NIPS*, J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta, Eds. Curran Associates, Inc., 2010, pp. 1759–1767.

[11] Y. Wu and P. Yang, "Minimax rates of entropy estimation on large alphabets via best polynomial approximation," *IEEE Trans. Inf. Theory*, vol. 62, no. 6, pp. 3702–3720, 2016.

[12] S. Fujishige, "Lexicographically optimal base of a polymatroid with respect to a weight vector," *Mathematics of Operations Research*, vol. 5, no. 2, pp. 186–196, 1980.

[13] ——, *Submodular functions and optimization*, 2nd ed. Elsevier, 2005.

[14] ——, "Theory of principal partitions revisited," in *Research Trends in Combinatorial Optimization*. Springer, 2009, pp. 127–162.

[15] C. Chan, A. Al-Bashabsheh, Q. Zhou, and T. Liu, "Duality between feature selection and data clustering," in *54th Annual Allerton Conference on Communication, Control, and Computing, Allerton Retreat Center, Monticello, Illinois*, 2016.

[16] P. Narayan and H. Tyagi, "Multiterminal secrecy by public discussion," *Foundations and Trends in Communications and Information Theory*, vol. 13, no. 2-3, pp. 129–275, 2016.

[17] S. Fujishige, "Polymatroidal dependence structure of a set of random variables," *Information and Control*, vol. 39, no. 1, pp. 55 – 72, 1978.

[18] C. E. Shannon, "A mathematical theory of communication," *The Bell System Technical Journal*, vol. 27, no. 3, pp. 379–423, July 1948.

[19] A. Schrijver, *Combinatorial Optimization: Polyhedra and Efficiency*. Springer, 2002.

[20] S. Fujishige and S. Isotani, "A submodular function minimization algorithm based on the minimum-norm base," *Pacific Journal of Optimization*, vol. 7, no. 1, pp. 3–17, 2011.
[21] S. Watanabe, "Information theoretical analysis of multivariate correlation," *IBM Journal of Research and Development*, vol. 4, no. 1, pp. 66–82, 1960.
[22] C. Chan and T. Liu, "Clustering of random variables by multivariate mutual information on Chow-Liu tree approximations," in *Fifty-Third Annual Allerton Conference on Communication, Control, and Computing*, Allerton Retreat Center, Monticello, Illinois, Sep. 2015.
[23] V. Kolmogorov, "A faster algorithm for computing the principal sequence of partitions of a graph," *Algorithmica*, vol. 56, no. 4, pp. 394–412, 2010.
[24] C. Chan, A. Al-Bashabsheh, Q. Zhou, N. Ding, T. Liu, and A. Sprintson, "Successive omniscience," *IEEE Trans. Inf. Theory*, vol. 62, no. 6, pp. 3270–3289, June 2016.